

A Comparison of Directional Distances for Hand Pose Estimation

Dimitrios Tzionas^{1,2} and Juergen Gall²

¹ Perceiving Systems Department, MPI for Intelligent Systems, Germany
dimitris.tzionas@tuebingen.mpg.de

² Computer Vision Group, University of Bonn, Germany
gall@informatik.uni-bonn.de

Abstract. Benchmarking methods for 3d hand tracking is still an open problem due to the difficulty of acquiring ground truth data. We introduce a new dataset and benchmarking protocol that is insensitive to the accumulative error of other protocols. To this end, we create testing frame pairs of increasing difficulty and measure the pose estimation error separately for each of them. This approach gives new insights and allows to accurately study the performance of each feature or method without employing a full tracking pipeline. Following this protocol, we evaluate various directional distances in the context of silhouette-based 3d hand tracking, expressed as special cases of a generalized Chamfer distance form. An appropriate parameter setup is proposed for each of them, and a comparative study reveals the best performing method in this context.

1 Introduction

Benchmarking methods for 3d hand tracking has been identified in the review [7] as an open problem due to the difficulty of acquiring ground truth data. As in one of the earliest works on markerless 3d hand tracking [19], quantitative evaluations are still mostly performed on synthetic data, e.g., [26,2,34,4,21]. The vast majority of the related literature, however, is limited to visual, qualitative performance evaluation, where the estimated model is overlaid on the images.

While there are several datasets and evaluation protocols for benchmarking human pose estimation methods publicly available, where markers [28,1], inertial sensors [3], or a semi-automatic annotation approach [32] have been used to acquire ground truth data, there are no datasets available for benchmarking articulated hand pose estimation. We propose thus a benchmark dataset consisting of 4 sequences of two interacting hands captured by 8 cameras, where the ground truth position of the 3d joints has been manually annotated.

Tracking approaches are usually evaluated by providing the pose for the first frame and measuring the accumulative pose estimation error for all consecutive frames of the sequence, e.g., [28]. While this protocol is optimal for comparing full tracking systems, it makes it difficult to analyze the impact of individual components of a system. For instance, a method that estimates the joint positions with a high accuracy, but fails in a few cases and is unable to recover from errors,




Fig. 1: *Initial pose* (a) and synthetic (b) and realistic (b) *target silhouettes* of one camera view. The benchmark measures the pose estimation error of the joints of both hands. In the synthetic experiments all joints (*all* dots in (a)) are taken into account, while in the realistic only a subset (*black* dots in (a)) is evaluated.

will have a high tracking error if an error occurs very early in a test sequence. However, the tracking error will be very low if the error occurs at the end of the sequence. The accumulation of tracking errors makes it difficult to analyze in-depth situations where an approach works or fails. We therefore propose a benchmark that analyzes the error not over a full sequence, but over a set of pairs consisting of a starting pose and a test frame. Based on the start pose and the test frame, the pairs have different grades of difficulty.

In this work, we use the proposed benchmark to analyze various silhouette-based distance measures for hand pose estimation. Distance measures that are based on a closest point distance, like the Chamfer distance, are commonly used due to its efficiency [19] and often extended by including directional information [9,33]. Recently, a fast method that computes a directional Chamfer distance using a 3d distance tensor has been proposed [16] for shape matching. In this work, we introduce a general form of the Chamfer distance for hand pose estimation and quantitatively compare several special cases.

2 Related Work

Since the earliest days of vision-based hand pose estimation [24,7], low-level features like silhouettes [19], edges [13], depth [6], optical flow [19], shading [14] or a combination of them [17] have been used for hand pose estimation. Although Chamfer distances combined with an edge orientation term have been used in [33,2,31,29], the different distances have not been thoroughly evaluated for hand pose estimation. While a KD-tree is used in [31] to compute a directional Chamfer distance, Liu et al. [16] recently proposed a distance transform approach to efficiently use a directional Chamfer distance for shape matching. Different methods of shape matching for pose estimation have been compared in the context of rigid objects [12] or articulated objects [22]. While previous work mainly considered to estimate the pose of a hand in isolation, recent works con-

sider more complicated scenarios where two hands interact with each other [21,4] or with objects [11,25,10,20,4].

3 Hand Pose Estimation

For evaluation, we use a publicly available hand model [4], consisting of a set of vertices, an underlying kinematic skeleton with 35 degrees of freedom (DOF) per hand, and skinning weights. The vertices and the joints of the skeleton are shown in Fig. 1. Each 3d vertex \mathbf{v} is associated to a bone j by the skinning weights $\alpha_{\mathbf{v},j}$, where $\sum_j \alpha_{\mathbf{v},j} = 1$. The articulated deformations of a skeleton are encoded by the vector θ that represents the rigid bone transformations $T_j(\theta)$, i.e., rotation and translation, by twists $\hat{\xi} \in se(3)$ [18,5]. Each twist-encoded rigid body transformation $\theta_j \hat{\xi}_j$ for a bone j can be converted into a homogeneous transformation matrix by the exponential map operator, i.e., $T_j(\theta) = \exp(\theta_j \hat{\xi}_j) \in SE(3)$. The mesh deformations based on the pose parameters θ are obtained by the linear blend skinning operator [15] using homogeneous coordinates:

$$\mathbf{v}(\theta) = \sum \alpha_{\mathbf{v},j} T_j(\theta) \mathbf{v} . \quad (1)$$

In order to estimate the hand pose for a given frame, correspondences between the mesh and the image of each camera c are established. Each correspondence $(\mathbf{v}_i, \mathbf{q}_i, c_i)$ associates a vertex \mathbf{v}_i to a 2d point \mathbf{q}_i in camera view c_i . Assuming that the cameras are calibrated, the point \mathbf{q}_i can be converted into a projection ray that is represented by the direction \mathbf{d}_i and moment \mathbf{m}_i of the line [30,27]. The hand pose can then be determined by the pose parameters that minimize the shortest distance between the 3d vertices \mathbf{v}_i and 3d projection rays $(\mathbf{d}_i, \mathbf{m}_i)$:

$$\operatorname{argmin}_{\theta} \frac{1}{2N} \sum_{i=1}^N \|\mathbf{v}_i(\theta) \times \mathbf{d}_i - \mathbf{m}_i\|^2 . \quad (2)$$

This non-linear least-squares problem can be iteratively solved [27]:

- Extract correspondences for all cameras $(\mathbf{v}_i, \mathbf{q}_i, c_i)$,
- Solve (2) using the linearization $T_j(\theta) = \exp(\theta_j \hat{\xi}_j) \approx I + \theta_j \hat{\xi}_j$,
- Update vertex positions by (1).

In this work, we reformulate (2) as a Chamfer distance minimization problem.

4 Generalized Chamfer Distance

As discussed in Section 2, the Chamfer distance is commonly used for shape matching and has been also used for pose estimation by shape matching. In our context, the Chamfer distance between pixels of a contour \mathcal{C} for a given camera view and the set of projected rim vertices $\mathcal{P}(\theta)$, which depend on the pose parameters θ and project onto the contour of the projected surface, is

$$d_{\text{Chamfer}}(\theta, \mathcal{C}) = \frac{1}{|\mathcal{P}(\theta)|} \sum_{\mathbf{p} \in \mathcal{P}(\theta)} \min_{\mathbf{q} \in \mathcal{C}} \|\mathbf{p} - \mathbf{q}\| . \quad (3)$$

This expression can be efficiently computed using a 2d distance transform [8].

The Chamfer distance (3) can be generalized by

$$d_{Chamfer}^{Z,f,d}(\theta, \mathcal{C}) = \frac{1}{Z} \sum_{\mathbf{p} \in \mathcal{P}(\theta)} f\left(\mathbf{p}, \underset{\mathbf{q} \in \mathcal{C}}{\operatorname{argmin}} d(\mathbf{p}, \mathbf{q})\right), \quad (4)$$

where $d(\mathbf{p}, \mathbf{q})$ is a 2d distance function to compute the distance between two points, $f(\mathbf{p}, \mathbf{q})$ is a penalty function for two closest points, and Z is a normalization factor. If we use

$$d(\mathbf{p}, \mathbf{q}) = \|\mathbf{p} - \mathbf{q}\|, \quad f(\mathbf{p}, \mathbf{q}) = d(\mathbf{p}, \mathbf{q}), \quad Z = |\mathcal{P}(\theta)|, \quad (5)$$

$d_{Chamfer}^{Z,f,d}(\theta, \mathcal{C})$ is the standard Chamfer distance (3). In order to increase the robustness to outliers, $f(\mathbf{p}, \mathbf{q}) = \min(d(\mathbf{p}, \mathbf{q})^2, K)$ is used in [29], where K is a threshold on the maximum squared distance.

Orientation can be integrated by penalizing correspondences with inconsistent orientations:

$$d(\mathbf{p}, \mathbf{q}) = \|\mathbf{p} - \mathbf{q}\|, \quad f(\mathbf{p}, \mathbf{q}) = \begin{cases} d(\mathbf{p}, \mathbf{q}) & \text{if } |\phi(\mathbf{p}) - \phi(\mathbf{q})|_\phi < \tau \\ K & \text{otherwise} \end{cases}, \quad Z = |\mathcal{P}(\theta)|, \quad (6)$$

or by computing the closest distance to points of similar orientation based on a circular distance threshold τ [9]:

$$d(\mathbf{p}, \mathbf{q}) = \begin{cases} \|\mathbf{p} - \mathbf{q}\| & \text{if } |\phi(\mathbf{p}) - \phi(\mathbf{q})|_\phi < \tau \\ \infty & \text{otherwise} \end{cases}, \quad f(\mathbf{p}, \mathbf{q}) = d(\mathbf{p}, \mathbf{q}), \quad Z = |\mathcal{P}(\theta)|, \quad (7)$$

where $|\phi(\mathbf{p}) - \phi(\mathbf{q})|_\phi$ is the circular distance between two angles, which can be signed, i.e., in the range of $[0, \pi]$, or unsigned, i.e., in the range of $[0, \frac{\pi}{2}]$.

The directional Chamfer distance [16] can be written as

$$d(\mathbf{p}, \mathbf{q}) = \|\mathbf{p} - \mathbf{q}\| + \lambda |\phi(\mathbf{p}) - \phi(\mathbf{q})|_\phi, \quad f(\mathbf{p}, \mathbf{q}) = d(\mathbf{p}, \mathbf{q}), \quad Z = |\mathcal{P}(\theta)|. \quad (8)$$

To compute $d_{Chamfer}^{Z,f,d}(\theta, \mathcal{C})$ with (8) efficiently, ϕ can be quantized to compute a 3d distance transform [16]. As in [16], we compute $\phi(\mathbf{q})$ by converting \mathcal{C} into a line representation [23]. $\phi(\mathbf{p})$ is obtained by projecting the normals of the corresponding vertices in $\mathcal{P}(\theta)$.

In order to use the generalized Chamfer distance $d_{Chamfer}^{Z,f,d}(\theta, \mathcal{C})$ for pose estimation from multiple views (2), only f and Z need to be adapted. Let $\mathcal{C}(c)$ denote the contour of camera view c and $\mathcal{P}(\theta, c)$ the set of projected vertices for pose parameters θ and camera c . (2) can be rewritten as

$$\underset{\theta}{\operatorname{argmin}} \frac{1}{2 \sum_c |\mathcal{P}(\theta, c)|} \sum_c d_{Chamfer}^{Z,f,d}(\theta, \mathcal{C}(c)) \quad (9)$$

$$\text{with} \quad f(\mathbf{p}, \mathbf{q}) = \|\mathbf{v}(\theta) \times \mathbf{d} - \mathbf{m}\|^2, \quad Z = 1, \quad (10)$$

where $\mathbf{v}(\theta)$ is the 3d vertex corresponding to $\mathbf{p} \in \mathcal{P}(\theta)$ and (\mathbf{d}, \mathbf{m}) is the 3d projection ray corresponding to \mathbf{q} . $d(\mathbf{p}, \mathbf{q})$ can be any of the functions (5)-(8).

In case of (6), instead of adding a fixed penalty term K , correspondences with inconsistent orientation can be simply removed and $\mathcal{P}(\theta, c)$ becomes the set of correspondences with $|\phi(\mathbf{p}) - \phi(\mathbf{q})|_\phi < \tau$.

5 Benchmark

We propose a benchmarking protocol that analyzes the error not over full sequences, but over a sampled *set of testing pairs*. Each pair consists of a *starting pose* and a *test frame*, ignoring the intermediate frames to simulate various difficulties. This approach gives new insights and provides means to analyze in-depth the contributions of various features or methods to the overall tracking pipeline under varying difficulty and to thoroughly study failure cases.

In this respect, 4 publicly available sequences¹ are used, containing realistic scenarios of two strongly interacting hands [4]. 10% of the total frames are randomly selected, forming the set of *test frames* of the final pairs. This is the basis to create 4 different sets of image pairs, having 1,5,10,15 frames difference respectively between the *starting pose* and the *test frame*, presenting thus increasing difficulty for tracking systems. These 4 sets and the overall combination constitute a challenging dataset, representing realistic scenarios the occur due to low frame rates, fast motion or estimation errors in the previous frame.

The created testing sets are used in two experimental setups: a purely *synthetic* and a *realistic*. In both cases, the *starting pose* is given by the publicly available motion data outputted by the tracker of [4]. In the *synthetic* experimental setup the *test frame* is synthesized by the hand model and the aforementioned motion data, while the required ground truth exists inherently in them. In the *realistic* setup the *test frame* is given by the camera images, for which no ground-truth data are available, thus the frames have been manually annotated². As error measure, we use the average of the Euclidean distances between the estimated and the ground-truth 3d positions of the joints. For the realistic setup we use only the joints of the model that could be annotated, which are depicted with black color in Fig. 1. For the synthetic setup all joints of the model (black and red) are taken into account.

6 Experiments

6.1 Implementation Details

The aforementioned benchmark is used to evaluate four special cases of the generalized Chamfer distance (Section 4) for hand pose estimation.

CH denotes the Chamfer distance without any orientation information (5).

DCH-Thres rejects correspondences if the orientations are inconsistent, depending on the circular distance threshold τ (6).

DCH-Quant computes a 2d distance field for all quantizations of ϕ and assigns a vertex to one bin based on the orientation of its normal (7). Instead of hard binning, soft binning can also be performed, denoted by **DCH-Quant2**. In this case, the two closest bins are used, yielding two correspondences per vertex.

¹ Model, videos, and motion data are provided at <http://cvg.ethz.ch/research/ih-mocap>. Sequences: *Finger tips touching and praying*, *Fingers crossing and twisting*, *Fingers folding*, *Fingers walking*. Video: 1080 × 1920 px, 50 fps, 8 camera-views.

² The ground-truth annotated dataset, along with a viewer-application, is available at http://files.is.tue.mpg.de/dtzionas/GCPR_2013.html.

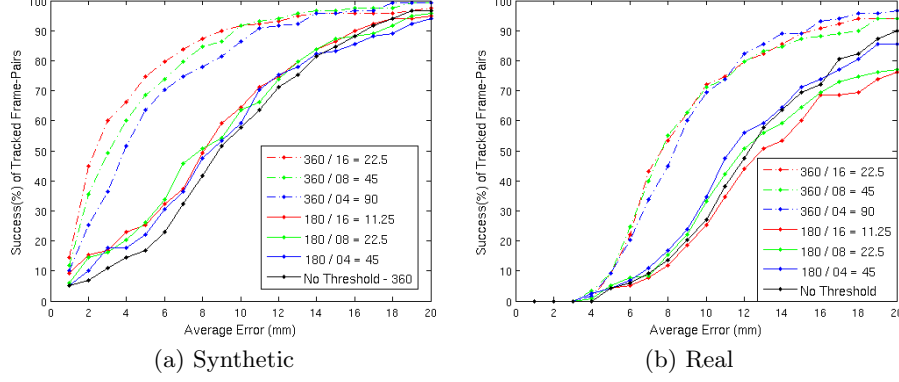


Fig. 2: Performance evaluation of **DCH-Thres** with different values of τ and both signed (360) and unsigned (180) distance $|\cdot|_\phi$. The plots show the percentage of frame pairs (y -axis) below a given average error (x -axis). The *signed* distance (360) significantly outperforms the *unsigned* distance (180), and the best performing circular distance threshold value is $\tau = 22.5$.

DCH-DT3 denotes the approximation of the directional Chamfer distance (8) proposed by Liu et al. [16]. The approach computes a 3d distance field $DT3$ and depends on two parameters. While λ steers the impact of the orientation term in (8), ϕ is quantized by a fixed number of bins.

As mentioned in Section 4, the *target silhouette* is approximated with linear line segments for all the directional distances DCH , using [23]. We also investigate two versions of the circular distance $|\cdot|_\phi$, namely the unsigned version, denoted by *180*, and the signed version, denoted by *360*.

6.2 Results

We have evaluated all Chamfer distances both on the synthetic and the realistic dataset in order to compare the distances for 3d hand pose estimation, but also in order to investigate the performance predicting abilities of synthetic test data. As measure, we use the average joint error per test frame and compute the percentage of frames with an error below a given threshold. We first evaluated the differences between the signed and unsigned circular distance for $DCH-Thres$ and varied the threshold parameter τ . The results are plotted in Fig. 2. The plot shows that the signed distance outperforms the unsigned distance. Since we observed the same result for $DCH-DT3$, we only report results for the signed distance (360) in the remaining experiments.

For $DCH-DT3$, we evaluated the impact of the two parameters λ and the number of quantization bins for ϕ . The results are plotted in Fig. 3. Figs. 3a and 3b show the importance of directional information for hand pose estimation, and reveal that there is a large range of λ that works well. With a finer quantization of ϕ , the original directional Chamfer distance (8) is better approximated. Figs. 3c and 3d show that 16 bins are sufficient for this task.

Table 1: Mean error \pm std.dev.(mm), av.time (sec) for 1,5,10,15 frame differences. Time measurements regard single-threaded code on a 6-core 3GHz Xeon PC.

		1	5	10	15	All	Time
Synthetic	<i>CH</i>	1.0 \pm 1.0	2.5 \pm 2.5	4.3 \pm 4.6	6.4 \pm 6.1	3.5 \pm 4.5	103
	<i>DCH-DT3</i>	2.0 \pm 1.3	2.3 \pm 1.3	3.8 \pm 2.9	6.2 \pm 5.8	3.6 \pm 3.8	115
	<i>DCH-Quant</i>	4.0 \pm 1.6	4.2 \pm 1.7	5.4 \pm 2.5	7.0 \pm 4.0	5.1 \pm 2.9	161
	<i>DCH-Thres</i>	1.1 \pm 0.8	1.3 \pm 1.1	2.5 \pm 2.4	4.1 \pm 4.5	2.2 \pm 2.9	077
Realistic	<i>Initial</i>	6.4 \pm 2.0	10.5 \pm 5.6	16.5 \pm 11.5	22.6 \pm 16.9	14.0 \pm 12.3	-
	<i>Ballan et al. [4]</i>	5.9 \pm 1.9	-	-	-	-	-
	<i>CH</i>	7.1 \pm 1.9	7.8 \pm 2.4	9.3 \pm 4.3	10.9 \pm 5.9	8.8 \pm 4.2	-
	<i>DCH-DT3</i>	6.3 \pm 1.5	6.7 \pm 2.0	8.7 \pm 5.1	11.1 \pm 7.9	8.3 \pm 5.4	-
	<i>DCH-Quant</i>	6.8 \pm 1.6	7.2 \pm 2.1	9.0 \pm 4.4	10.7 \pm 7.3	8.4 \pm 4.7	-
	<i>DCH-Thres</i>	6.1 \pm 1.3	6.4 \pm 1.8	7.6 \pm 3.3	9.4 \pm 5.3	7.4 \pm 3.6	-

We finally evaluated the number of bins for *DCH-Quant* and *DCH-Quant2*. Fig. 4 shows that *DCH-Quant2* performs better than *DCH-Quant*. In this case, a large number of bins results in a very orientation sensitive measure, and the performance decreases with a finer quantization, in contrast to *DCH-DT3*.

Fig. 5 summarizes the results for each distance with the best parameter setting. As expected, the results show that directional information improves the estimation accuracy. However, it is not *DCH-DT3* that performs best for hand pose estimation, but *DCH-Thres*, which is also more efficient to compute. While for *DCH-DT3* the full hand model converges smoothly to the final pose, the thresholding yields a better fit to the silhouette after convergence (*see supplementary video*³). Comparing the performances between synthetic and real data, we conclude that synthetic data is a good performance indicator, but might be misleading sometimes. For instance, *CH* performs well on the synthetic data but worst on the real data. This is also reflected by the mean error for the various frame differences provided in Table 1, that introduce an increasing difficulty in the benchmark. Denoted with the term *initial* is the average 3d distance of the joints before running the pose estimation algorithm. The result of a full tracking system [4] is provided for comparison, which expectedly performs better due to the number of features combined. Finally, runtime is provided for the synthetic experiments, giving some intuition about the time efficiency of each method.

7 Conclusion

In this work, we propose a new benchmark dataset for hand pose estimation that allows to evaluate single components of a hand tracker without running a full system. As an example, we discuss a generalized Chamfer distance and evaluate four special cases. The experiments reveal that directional information is important and a signed circular distance performs better than an unsigned distance in the case of silhouettes. Interestingly, a distance using a circular threshold outperforms a smooth directional Chamfer distance both in terms of accuracy and

³ <http://youtu.be/Cbu3eEc11qk>

runtime. We finally conclude that synthetic data can be a good indicator for the performance, but might be misleading when comparing different methods. Future plans include adding frame pairs of other sequences with more background clutter and segmentation noise.

8 Acknowledgments

The authors acknowledge financial support from the DFG Emmy Noether program (GA 1927/1-1) and the Max Planck Society.

References

1. van der Aa, N.P., Luo, X., Giezeman, G.J., Tan, R.T., Veltkamp, R.C.: Umpm benchmark: A multi-person dataset with synchronized video and motion capture data for evaluation of articulated human motion and interaction. In: Workshop on Human Interaction in Computer Vision. pp. 1264–1269 (2011)
2. Athitsos, V., Sclaroff, S.: Estimating 3d hand pose from a cluttered image. In: CVPR. pp. 432–439 (2003)
3. Baak, A., Helten, T., Müller, M., Pons-Moll, G., Rosenhahn, B., Seidel, H.P.: Analyzing and evaluating markerless motion tracking using inertial sensors. In: Workshop on Human Motion. pp. 137–150 (2010)
4. Ballan, L., Taneja, A., Gall, J., Van Gool, L., Pollefeys, M.: Motion capture of hands in action using discriminative salient points. In: ECCV. pp. 640–653 (2012)
5. Bregler, C., Malik, J., Pullen, K.: Twist based acquisition and tracking of animal and human kinematics. IJCV 56(3), 179–194 (2004)
6. Delamarre, Q., Faugeras, O.D.: 3d articulated models and multiview tracking with physical forces. CVIU 81(3), 328–357 (2001)
7. Erol, A., Bebis, G., Nicolescu, M., Boyle, R.D., Twombly, X.: Vision-based hand pose estimation: A review. CVIU 108(1-2), 52–73 (2007)
8. Felzenszwalb, P.F., Huttenlocher, D.P.: Distance transforms of sampled functions. Theory of Computing 8(19), 415–428 (2012)
9. Gavrilu, D.: Multi-feature hierarchical template matching using distance transforms. In: ICPR. pp. 439–444 (1998)
10. Hamer, H., Gall, J., Weise, T., Van Gool, L.: An object-dependent hand pose prior from sparse training data. In: CVPR. pp. 671–678 (2010)
11. Hamer, H., Schindler, K., Koller-Meier, E., Van Gool, L.: Tracking a hand manipulating an object. In: ICCV. pp. 1475–1482 (2009)
12. Han, D., Rosenhahn, B., Weickert, J., Seidel, H.P.: Combined registration methods for pose estimation. In: ISVC. pp. 913–924 (2008)
13. Heap, T., Hogg, D.: Towards 3d hand tracking using a deformable model. In: FG. pp. 140–145 (1996)
14. de La Gorce, M., Fleet, D.J., Paragios, N.: Model-based 3d hand pose estimation from monocular video. PAMI 33(9), 1793–1805 (2011)
15. Lewis, J.P., Corder, M., Fong, N.: Pose space deformation: a unified approach to shape interpolation and skeleton-driven deformation. In: SIGGRAPH (2000)
16. Liu, M.Y., Tuzel, O., Veeraraghavan, A., Chellappa, R.: Fast directional chamfer matching. In: CVPR. pp. 1696–1703 (2010)
17. Lu, S., Metaxas, D., Samaras, D., Oliensis, J.: Using multiple cues for hand tracking and model refinement. In: CVPR. pp. 443–450 (2003)

18. Murray, R.M., Sastry, S.S., Zexiang, L.: A Mathematical Introduction to Robotic Manipulation. CRC Press, Inc., Boca Raton, FL, USA (1994)
19. Nirei, K., Saito, H., Mochimaru, M., Ozawa, S.: Human hand tracking from binocular image sequences. In: IECON. pp. 297–302 (1996)
20. Oikonomidis, I., Kyriazis, N., Argyros, A.: Full dof tracking of a hand interacting with an object by modeling occlusions and physical constraints. In: ICCV (2011)
21. Oikonomidis, I., Kyriazis, N., Argyros, A.A.: Tracking the articulated motion of two strongly interacting hands. In: CVPR. pp. 1862–1869 (2012)
22. Pons-Moll, G., Leal-Taixé, L., Truong, T., Rosenhahn, B.: Efficient and robust shape matching for model based human motion capture. In: DAGM (2011)
23. Ramer, U.: An iterative procedure for the polygonal approximation of plane curves. *Computer Graphics and Image Processing* 1(3), 244 – 256 (1972)
24. Rehg, J.M., Kanade, T.: Visual tracking of high dof articulated structures: an application to human hand tracking. In: ECCV. pp. 35–46 (1994)
25. Romero, J., Kjellström, H., Kragic, D.: Hands in action: real-time 3d reconstruction of hands in interaction with objects. In: ICRA. pp. 458–463 (2010)
26. Rosales, R., Athitsos, V., Sigal, L., Sclaroff, S.: 3d hand pose reconstruction using specialized mappings. In: ICCV. pp. 378–387 (2001)
27. Rosenhahn, B., Brox, T., Weickert, J.: Three-dimensional shape knowledge for joint image segmentation and pose tracking. *IJCV* 73, 243–262 (2007)
28. Sigal, L., Balan, A., Black, M.: Humaneva: Synchronized video and motion capture dataset and baseline algorithm for evaluation of articulated humanmotion. *IJCV* 87, 4–27 (2010)
29. Stenger, B., Thayananthan, A., Torr, P.: Model-based hand tracking using a hierarchical bayesian filter. *PAMI* 28(9), 1372–1384 (2006)
30. Stolfi, J.: Oriented Proj. Geometry: A Framework for Geom. Computation. Academic Press, Boston (1991)
31. Sudderth, E., Mandel, M., Freeman, W., Willsky, A.: Visual Hand Tracking Using Nonparametric Belief Propagation. In: Workshop on Generative Model Based Vision. pp. 189–189 (2004)
32. Tenorth, M., Bandouch, J., Beetz, M.: The TUM Kitchen Data Set of Everyday Manipulation Activities for Motion Tracking and Action Recognition. In: Int.Work. on Tracking Humans for the Eval. of their Motion in Im.Seq. pp. 1089–1096 (2009)
33. Thayananthan, A., Stenger, B., Torr, P.H.S., Cipolla, R.: Shape context and chamfer matching in cluttered scenes. In: CVPR. pp. 127–133 (2003)
34. Zhou, H., Huang, T.: Okapi-chamfer matching for articulate object recognition. In: ICCV. pp. 1026–1033 (2005)




Fig. 3: **(a-b)** Performance evaluation of **DCH-DT3** with different values of λ , using 16 quantization bins. While the orientation term significantly improves the performance, the performance gets saturated for values in the range [15,35]. **(c-d)** Performance evaluation of **DCH-DT3** with different quantizations of ϕ , using $\lambda = 25$. The synthetic data shows that more than 8 bins are required, though the differences are rather small on the real dataset. This is in accordance with Fig. 2 since a threshold of 22.5 corresponds to 16 quantization bins.

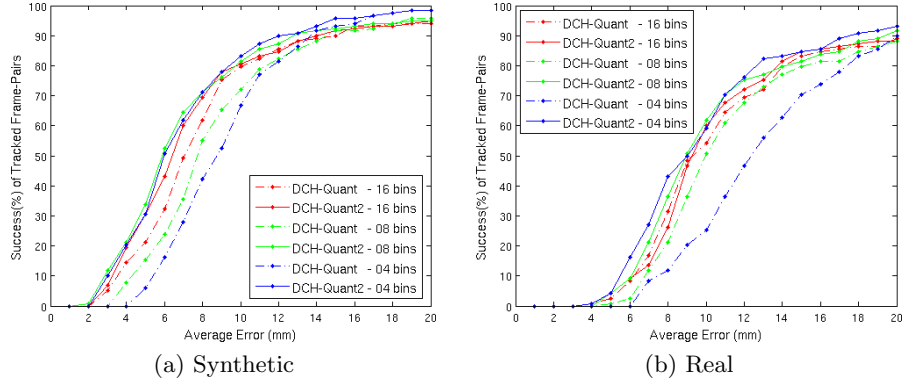


Fig. 4: Performance evaluation of **DCH-Quant** and **DCH-Quant2** with different quantizations of ϕ . Soft-binning outperforms hard assignments and in this case fewer bins perform better than many bins.




Fig. 5: Comparison of all distances with best settings. Although *DCH-DT3* provides a smoother distance measure, *DCH-Thres* performs best on both datasets.